Decision Ownership Is The Missing Layer In AI Governance

By Jay Hawkinson | Forbes Technology Council Member

Originally published: Forbes Technology Council, April 13, 2026

One of the most common ways organizations struggle with AI has nothing to do with the technology itself. It comes down to decisions. Across industries, companies are deploying models quickly and focusing on performance, accuracy and speed. But in many cases, they are not asking a fundamental question: Who is accountable for the decisions those systems influence?

Without a clear answer, organizations end up in a dangerous position. Models generate outputs and teams act on them, but when something goes wrong, no one owns the consequence. That gap can lead to costly mistakes, reputational damage and regulatory exposure, regardless of the size of the business. The root of the problem lies in how organizations define ownership of decisions in the first place.

Decision Ownership: An Organizational Design Problem

The first mistake companies make is treating decision ownership as a data governance question. In fact, it is an organizational design question, one that has to be answered before any model goes into production.

What does decision ownership actually mean in practice? When the AI output is wrong, someone's name is attached to the consequence, and they have both the authority to act and the accountability for the result.

The way to establish that accountability is to map every AI-influenced decision just like you would map any high-stakes business process. Who was authorized to make this decision before AI was involved? What does AI change about the speed, the information available or the range of options? And critically: Does the decision now move faster than your approval structure can absorb? If the answer to that last question is "yes," you have an authority gap, and governance alone will not fix it.

Building Governance That Accelerates

The framing of governance versus innovation is the wrong starting point. Unclear governance is what slows innovation because teams spend their time relitigating questions of authority that should have been resolved before the project started.

The governance model that actually works at scale has three layers, and they need to be built in order. The foundation is what I call authoritative architecture: who owns which decisions, what the escalation structure looks like and where human review is required. This is an operating model specification that tells every team, before deployment, where the boundaries are and who to call when something goes wrong.

The second layer is accountability mapping. For every AI-influenced decision, there is a named owner who has the authority to act and accepts the consequences. If the person accountable for an outcome lacks the authority to change the system producing it, the governance structure is incomplete.

The third layer is escalation logic: specific, predefined triggers that route decisions back to human review. This means concrete thresholds, not just general policies. Human review kicks in automatically when model confidence falls below a defined level, when a decision involves an amount above a certain threshold or when output contradicts a prior decision made by a senior leader. Teams that define these triggers in advance actually move faster because they are not relitigating authority in real time.

For data and analytics leaders, the practical starting point is an audit of your current decision architecture with no reference to AI at all. Map the 10 most consequential decisions in your organization. For each one, ask: Who has the authority to make it, what information is required and what is the escalation path if it goes wrong? Most organizations find that several of those decisions are made through informal authority structures that have never been documented. AI is great at exposing that problem and accelerating it.

Bring the authority conversation to the board before the board comes to you. Boards are being asked to exercise AI oversight without a clear picture of where authority sits in the operating model below them. Data leaders who can present a clear decision authority map, with named owners and defined escalation paths, give the board what it truly needs.

The organizations that scale AI effectively over the next five years will be those that understand the purpose behind the decisions before they ever deploy the technology. We have made this mistake before with other technologies, spending billions without asking the most fundamental question: Do we really understand what decisions we are trying to make, and who is responsible for making them?

ABOUT THE AUTHOR

Jay Hawkinson is a board-certified (NACD.DC) data and digital leader turning data into margin through AI, analytics and strategy. Founder, Hawksroost LLC. Contact: jhawkinson@hawksroost.com | hawksroost.com

This article was originally published on Forbes Technology Council on April 13, 2026. The author retains copyright. Republished with attribution per Forbes Councils member guidelines. Original: forbes.com/councils/forbestechcouncil/people/chrisjayhawkinson/